

Open-Set Vein Biometric Recognition with Deep Metric Learning

Paweł Pilarek¹[0009-0003-5955-9909]*,
Marcel Musiałek¹[0009-0009-4964-0547]*, and
Anna Górska¹[0000-0003-3305-8711]

¹Wrocław University of Science and Technology, Wrocław, Poland
278796@student.pwr.edu.pl, 279704@student.pwr.edu.pl,
anna.gorska@pwr.edu.pl

Abstract. Most state-of-the-art vein recognition methods rely on closed-set classification, which inherently limits their scalability and prevents the adaptive enrollment of new users without complete model retraining. We rigorously evaluate the computational boundaries of Deep Metric Learning (DML) under strict open-set constraints. Unlike standard closed-set approaches, we analyze the impact of data scarcity and domain shift on recognition performance. Our approach learns discriminative ℓ_2 -normalised embeddings and employs prototype-based matching with a calibrated similarity threshold to effectively distinguish between enrolled users and unseen impostors. We evaluate the framework under a strict subject-disjoint protocol across four diverse datasets covering finger, wrist, and dorsal hand veins (MMCBNU_6000, UTFVP, FYO, and a dorsal hand-vein dataset). On the large-scale MMCBNU_6000 benchmark, our best model (ResNet50-CBAM) achieves an OSCR of 0.9945, AUROC of 0.9974, and EER of 1.57%, maintaining high identification accuracy (99.6% Rank-1) while robustly rejecting unknown subjects. Cross-dataset experiments evaluate the framework’s generalisation across different acquisition setups, confirming that while the model handles large-scale data robustly, performance remains sensitive to domain shifts in low-data regimes. Ablation studies demonstrate that triplet-based objectives combined with a simple 1-NN classifier offer an optimal trade-off between accuracy and efficiency, enabling real-time deployment on commodity hardware.

Keywords: Vein Biometrics · Open-Set Recognition · Deep Metric Learning · Domain Shift Analysis · Computational Biometrics.

1 Introduction

Biometric recognition has become a cornerstone of modern access control, offering higher security than password-based mechanisms. While adopted modalities

* Equal contribution

such as fingerprint, iris, and voice recognition established significant benchmarks [8, 16], they face limitations: surface-level traits are susceptible to abrasion and forgery, while face and voice struggle with environmental noise and presentation attacks.

In this context, vascular (vein) biometrics offer a robust alternative due to high spoofing resistance. Located beneath the skin and captured via near-infrared (NIR) imaging, vein patterns are difficult to replicate and leave no latent traces [7, 2]. Furthermore, vascular structures are considered stable over time and exhibit high distinctiveness.

Vein biometrics span finger, palm, dorsal, and wrist modalities. While early methods used texture descriptors, deep learning has recently become the dominant paradigm [6, 9], with convolutional neural networks (CNNs) achieving substantial gains in segmentation and matching. However, the vast majority of these methods are evaluated under a closed-set assumption [14, 20]. In this scenario, all test identities are known beforehand, which fails to reflect realistic deployments where unseen subjects appear. A reliable system must not only recognize enrolled users but also reject non-enrolled (unknown) identities.

This aligns with open-set recognition, where models handle unknown classes absent during training [5, 12]. While extensively studied in generic vision using strategies like calibrated scores or margin-based losses, open-set capabilities remain critical for security against unknown adversaries [1, 3]. In contrast, open-set recognition for vein biometrics remains underexplored, with few works investigating specific architectures or loss functions [4, 1]. This leaves key questions: how well can deep metric learning support open-set scenarios, and how should decision rules be optimized for effective rejection?

In this work, we address these questions by investigating open-set vein recognition within a deep metric learning framework. We learn an embedding function mapping subjects to a feature space and define a decision rule combining similarity scores with a calibrated rejection threshold. We further compare metric learning losses, analyzing their impact on the balance between recognition accuracy and rejection capability.

The proposed framework is applicable to various vein types (finger, wrist, dorsal) and is independent of specific acquisition setups. Validated on public datasets using realistic protocols, our experimental results (evaluated via EER and AUC) demonstrate that the learned embeddings maintain high recognition performance for enrolled users while providing effective rejection of unknown identities, indicating that deep metric learning is a promising direction for practical open-set vein biometric systems.

2 Related Work

2.1 Vein Biometrics and Computational Challenges

Vein biometrics utilize internal vascular patterns captured via near-infrared (NIR) imaging, offering inherent resistance to spoofing and high privacy preser-

vation [7, 2]. While traditional hand-crafted pipelines based on texture or topology descriptors (e.g., maximum curvature, Gabor filters, LBP) established early benchmarks [7, 2], recent surveys indicate that their performance degrades significantly under strong intra-class variability and sensor changes. This computational limitation in modeling complex, non-linear variations necessitates the shift towards learned feature representations capable of handling open-set scenarios.

Concurrently, research has expanded into secure deployability. For instance, Ren et al. [15] proposed a finger-vein system processing RSA-encrypted images directly via a CNN. This underscores the necessity of balancing recognition accuracy with template protection—a key computational constraint in practical authentication systems which guides our proposed framework.

2.2 Deep Learning for Vein Recognition

Deep learning has reshaped the landscape of vein recognition over the last decade. Early work explicitly demonstrated that relatively shallow CNNs trained on finger-vein images can outperform traditional hand-crafted features [14]. Radzi et al. reported high identification rates on small-scale datasets using a four-layer CNN, illustrating the potential of learned features even under limited data conditions. Recent surveys by Hemis et al. and Kocakulak et al. distinguish biometric methods based on the task—specifically classification versus verification—and the underlying neural architecture, covering classical CNNs, Residual Networks (ResNet) and Attention-based models [6, 9]. These reviews conclude that such deep approaches consistently outperform traditional methods when combined with appropriate regularisation.

Several works have proposed specialised architectures tailored to vein patterns. Zhang et al. [20] introduced AGCNN, a network with learnable receptive fields and Gabor-like filters designed to capture line-like vein structures with reduced parameter complexity. Hybrid and multimodal systems have also been explored; for instance, Abdullahi et al. [1] presented a network fusing fingerprint and finger-vein data using temporal modelling to exploit sequence information. Generative models have also found application in this domain, particularly for addressing data scarcity. Zhang et al. [19] utilized Generative Adversarial Networks (GANs) to synthesise realistic vein images, a strategy that supports training in data-limited or open-set scenarios.

Deep learning has also been extended beyond finger veins. Marattukalam et al. [13] proposed a lightweight deep learning system for contactless wrist vascular recognition, demonstrating efficient operation on low-cost hardware. Furthermore, the template-protection scheme by Ren et al. [15] illustrates how deep architectures can be integrated with cryptographic mechanisms (e.g., RSA-based encryption) to ensure security without compromising recognition accuracy.

Despite this progress, most deep learning-based vein recognition methods are still developed and evaluated under closed-set verification or identification protocols. Models are typically trained on a fixed set of subjects and tested on disjoint images of the same subjects, with performance reported in terms of standard error rates [7, 6]. Only a small number of works explicitly examine

generalisation to unseen users [4], and even fewer design decision rules specifically for open-set operation, leaving open the question of how well deep vein systems can distinguish between enrolled and non-enrolled identities.

2.3 Open-Set Recognition in Biometrics

Open-set recognition (OSR) addresses the problem of correctly classifying samples from known classes while explicitly handling samples from unknown classes that were not observed during training. Comprehensive surveys on OSR, out-of-distribution detection and open-world recognition categorise existing methods into discriminative, generative and hybrid approaches, and discuss evaluation protocols and applications across computer vision and pattern recognition [5, 12, 3]. Common strategies include calibrating the outputs of closed-set classifiers, learning margin-based embeddings with compact intra-class clusters and large inter-class margins, modelling feature distributions and designing dedicated unknown detectors or rejection mechanisms.

In the biometric domain, OSR has been studied mainly for face recognition, but also for other modalities such as iris, gait and multimodal systems. Margin-based softmax variants and deep metric learning losses (e.g., ArcFace-style objectives) have been shown to improve open-set performance by enforcing better separation between genuine and impostor samples in the embedding space, often combined with feature-norm regularisation or score normalisation to stabilise decision thresholds [5, 4]. More recent work proposes open-set biometric frameworks that jointly optimise closed-set recognition and unknown-user rejection, and evaluate models on protocols that more closely resemble operational conditions across multiple modalities [3].

In contrast, open-set recognition for vein biometrics remains comparatively underexplored. Existing deep vein systems, including those based on CNNs, GAN-based augmentation, specialised architectures such as AGCNN, multi-modal fusion networks and wrist-vein models, are typically assessed on closed-set protocols and do not explicitly consider unknown-user rejection [14, 19, 20, 13, 1, 15]. Although margin-based losses and deep metric learning have established strong open-set baselines in other modalities such as face and iris recognition [5], to our knowledge, there is no systematic study that evaluates these frameworks for vein recognition under explicitly open-set protocols that separate training, enrolment, and evaluation identities across multiple vein types (e.g., finger, wrist, dorsal). This gap motivates the present work, in which we bring together advances in deep metric learning and open-set recognition and investigate how embedding design, loss functions, and calibrated decision rules affect open-set performance in vein-based biometric systems.

3 Proposed Method

In this section, we describe our proposed framework for open-set vein recognition. The system is designed to learn discriminative vascular representations from a

set of known identities and subsequently use these representations to identify enrolled users while rejecting unknown subjects. The pipeline consists of three main stages: (1) data partitioning and preprocessing, (2) deep metric learning for embedding extraction, and (3) open-set enrollment and inference.




Fig. 1. Overview of the proposed open-set vein recognition framework. The input vein image is mapped to a compact, ℓ_2 -normalised embedding by the ResNet50-CBAM backbone. During inference, the system computes cosine similarity scores against enrolled class prototypes. If the maximum similarity S_{max} falls below the calibrated threshold τ , the sample is rejected as an unknown identity; otherwise, it is assigned to the nearest enrolled class.

3.1 Open-Set Data Protocol and Preprocessing

To simulate realistic open-set conditions, we enforce a strict subject-disjoint protocol. Given one or more public vein datasets (e.g., MMCBNU_6000 and UTFVP for finger veins, FYO for wrist veins, and a dorsal hand-vein dataset), we treat each unique finger, wrist or hand instance as a distinct class. For each dataset, the subjects are first partitioned into two disjoint subsets: a known set used for training and enrolment, and an unknown set used exclusively to evaluate rejection capabilities. This ensures that the model never encounters the unknown identities during training. Concretely, for MMCBNU_6000 we use 60% / 10% / 30% of the known identities for training, validation and testing, respectively. For UTFVP, the corresponding proportions are 50% / 15% / 35%, while for FYO and the dorsal hand-vein dataset we use 55% / 10% / 35% splits. For evaluation, we adopt either a session-based or a sample-based protocol, depending on the dataset annotation. For datasets with explicit acquisition sessions, we use a session-based protocol in which enrolment images are taken from one session and probe images from different sessions. For datasets without session labels, we fall back to a sample-based protocol: for each enrolled class, a fixed number of images (e.g., 7 samples) are reserved for enrolment (gallery), while the remaining images (e.g., 3 samples) serve as probes for identification. When multiple datasets are used, we construct dataset-specific splits and optionally evaluate cross-dataset generalisation by using one database for training/enrolment and another as an open-set test source [7, 2].

During training, we employ a balanced $P \times K$ sampling strategy to construct mini-batches, where each batch contains P distinct identities and K images per identity, providing sufficient positive and negative pairs for metric learning. Basic image pre-processing and data augmentation follow the experimental setup described in Section 4.2.

3.2 Deep Embedding Framework

Backbone and Manifold Constraints. We utilize a ResNet50 backbone initialized on ImageNet and equipped with Convolutional Block Attention Modules (CBAM) to extract fine-grained vascular patterns. The network projects input ROIs into a compact embedding space ($D = 512$). Crucially, we enforce an ℓ_2 constraint on the output vectors, $\|\mathbf{z}\|_2 = 1$, restricting the feature manifold to a unit hypersphere. This normalization ensures that the optimization dynamics are driven purely by angular separability rather than magnitude variations inherent in NIR imaging, effectively stabilizing the training landscape.

Metric Learning Objective. To structure this manifold for open-set discrimination, we employ the Batch-Hard Triplet Loss with a margin $m = 0.3$. Unlike standard classification objectives, this loss operates directly on Euclidean distances. By utilizing a balanced $P \times K$ mini-batch sampler (where a *class* corresponds to a specific anatomical instance, e.g., a unique finger), the objective dynamically mines the hardest positive and negative pairs. This explicitly minimizes intra-class variance while enforcing a fixed margin between disjoint identities, which is computationally essential for robust open-set rejection.

Prototype-Based Inference. Following training, the system transitions to a prototype-based inference mode. For each enrolled identity i , we compute a class prototype \mathbf{c}_i as the ℓ_2 -normalized centroid of its enrollment embeddings. This representation allows for efficient similarity-based matching and naturally supports incremental enrollment without the need for backbone retraining.

3.3 Open-Set Decision Rule

For open-set recognition, the system must decide whether a query sample belongs to one of the enrolled identities or should be rejected as unknown. Given a query image \mathbf{q} , we compute its embedding $\mathbf{z}_q = f(\mathbf{q})$ and the cosine similarity scores to all class prototypes:

$$s_i = \mathbf{c}_i^\top \mathbf{z}_q, \quad i = 1, \dots, C,$$

where C is the number of enrolled identities. Let

$$S_{\max} = \max_i s_i, \quad \hat{i} = \arg \max_i s_i.$$

The final prediction is given by a thresholded maximum-similarity rule:

$$\text{prediction} = \begin{cases} \hat{i}, & \text{if } S_{\max} \geq \tau, \\ \text{Reject (Unknown)}, & \text{if } S_{\max} < \tau, \end{cases} \quad (1)$$

where τ is a decision threshold.

To determine an operating threshold τ , we use a validation set containing both known identities (disjoint from the enrollment set) and pseudo-unknowns (additional held-out known identities that are not enrolled). We sweep a range of candidate thresholds and select the value that maximises an open-set criterion (e.g., the Open-Set Classification Rate (OSCR) or AUROC) or satisfies a prescribed operating point such as a maximum tolerated false-positive rate on unknowns. This procedure is consistent with evaluation practices in the open-set recognition literature [5, 3].

3.4 Evaluation Metrics

We assess the system using both classical biometric metrics and open-set recognition metrics. For verification-oriented evaluation on the known subset, we report the Equal Error Rate (EER) and the True Positive Rate (TPR) at fixed False Positive Rate (FPR) operating points (e.g., 1% and 0.1% FPR), which are standard in vein biometrics [7, 2].

To quantify open-set performance, we treat the problem as a joint identification-and-detection task and report: (i) the Area Under the Receiver Operating Characteristic curve (AUROC) for distinguishing known versus unknown samples based on S_{\max} , and (ii) the Open-Set Classification Rate (OSCR) curve [5]. The OSCR curve plots the Correct Classification Rate (CCR) of known samples against the FPR on unknown samples as the decision threshold varies, providing a more informative view of the trade-off between accurate identification of enrolled users and rejection of impostors. Where appropriate, we summarise the curve either by its area (AUOSCR) or by CCR at specific FPR operating points, following recent open-set evaluation protocols[3].

4 Experiments

In this section, we detail the experimental setup, including the datasets, training configuration and evaluation protocols used to validate our open-set vein recognition framework. We then present quantitative results and ablation studies that analyse the impact of key components such as loss functions, embedding dimension and decision rules.

4.1 Datasets and Splits

We evaluate our method on four public vein biometric datasets: MMCBNU_6000 and UTFVP for finger veins, FYO for wrist veins and a dorsal hand-vein dataset.

- **MMCBNU_6000 Finger Vein:** This dataset contains finger-vein images from 100 subjects, with six fingers per subject (index, middle and ring fingers of both hands), resulting in 600 unique classes. Each class typically has around 10 samples.
- **UTFVP Finger Vein:** A widely used finger-vein benchmark with 60 subjects and 360 classes, containing multiple images per finger captured under controlled near-infrared illumination. The dataset provides 1,440 images in total.
- **FYO Wrist Vein:** A contactless wrist-vein dataset acquired with an NIR system, where each wrist is treated as a separate class. The dataset comprises 160 subjects with 160 classes and 1,920 images.
- **Dorsal Hand Vein:** A dorsal hand-vein dataset with 251 subjects and 502 classes (left and right hands treated as distinct classes), containing 1,782 images.

Table 1 summarises the basic statistics of the datasets used in our experiments, including the number of subjects, classes and images for each modality.

Table 1. Statistics of the vein datasets used in the experiments.

Dataset	#Subjects	#Classes	#Images	Modality
MMCBNU_6000	100	600	12,000	Finger vein
UTFVP	60	360	1,440	Finger vein
FYO	160	160	1,920	Wrist vein
Dorsal	251	502	1,782	Dorsal hand vein

To simulate a realistic open-set scenario, we enforce a strict subject-disjoint protocol. For each dataset, the total set of subjects is partitioned into two disjoint subsets: a *known* set (used for training and enrollment) and an *unknown* set (used exclusively for testing rejection). In our experiments, we typically use 70% of subjects as known and 30% as unknown. Within the known classes, we employ a session-based split: for each class, a fixed number of samples (e.g., 7 images) are reserved for enrollment (gallery), while the remaining samples (e.g., 3 images) are used as probes for identification. This ensures that the model is evaluated on unseen images of known subjects as well as completely unseen identities.

4.2 Training Configuration

The embedding network is trained using a balanced batch ($P \times K$) sampler. In each training iteration, we randomly sample $P = 16$ identities and $K = 4$ images per identity, creating a mini-batch of size $P \cdot K = 64$. This strategy ensures that every batch contains sufficient positive and negative pairs for effective metric learning.

Input images are preprocessed by converting grayscale samples to three-channel RGB (to ensure compatibility with standard backbones), resizing them

to a fixed resolution of 224×224 , and applying standard ImageNet normalization. To mitigate overfitting, we employ a streamlined data augmentation strategy during training consisting of random horizontal flipping (with probability $p = 0.3$), and apply a dropout rate of 0.5 in the embedding head. We employ a ResNet50 backbone equipped with CBAM modules as our default architecture, natively producing 512-dimensional ℓ_2 -normalised embeddings.

The model is optimized using the AdamW optimizer with a low initial learning rate of 3×10^{-5} (for stability under metric learning objectives) and an aggressive weight decay of 10^{-3} . A cosine-annealing scheduler is used to adjust the learning rate over 100 epochs. To prevent overfitting, we employ early stopping with a patience of 15 epochs, monitoring the validation loss on a held-out subset of the known classes ($\Delta_{min} = 10^{-4}$).

All experiments were conducted on a workstation equipped with a single NVIDIA GeForce RTX 4070 Ti SUPER GPU. The average training time for the ResNet50 backbone on the MMCBNU_6000 dataset was approximately 30 minutes.

4.3 Evaluation Protocol

Our open-set evaluation protocol consists of three phases: enrollment, calibration, and testing.

1. **Enrollment:** For each known class in the test split, we compute a class prototype by averaging the embeddings of the 7 enrollment samples. These prototypes form the gallery.
2. **Calibration:** We use a separate validation set (containing both known and pseudo-unknown samples) to calibrate the open-set decision threshold. We sweep over candidate threshold values and select the one that maximises the Open-Set Classification Rate (OSCR) or AUROC, balancing the trade-off between correct classification of knowns and rejection of unknowns.
3. **Testing:** During inference, query samples from both the known (probe) and unknown sets are compared against the enrolled prototypes. We employ a 1-Nearest Neighbour (1-NN) rule (or, in ablation studies, a k -NN variant) to determine the maximum similarity score. If the score exceeds the calibrated threshold, the sample is assigned to the corresponding class; otherwise, it is rejected as unknown.

4.4 Metrics

We report a comprehensive set of metrics to assess performance:

- **Detection Metrics:** AUROC (Area Under the ROC Curve) and EER (Equal Error Rate) measure the system’s ability to distinguish between known and unknown subjects based on similarity scores. These metrics are widely established as standard benchmarks for open-set recognition in biometrics literature [5, 4].

- **Open-Set Metrics:** The OSCR curve plots the Correct Classification Rate (CCR) against the False Positive Rate (FPR). We also report the True Positive Rate (TPR) at strict FPR targets, e.g. $\text{TPR@FPR} = 10^{-2}$ and 10^{-3} .
- **Identification Metrics:** For known subjects, we report Rank-1 and Rank-5 identification accuracy (CMC ranks).
- **Operational Metrics:** We report the final calibrated threshold, the accuracy on known subjects and the rejection rate on unknown subjects, as well as average runtime for enrollment and query inference.

All reported results are averaged over three random splits of the subjects into known and unknown sets, unless stated otherwise.

4.5 Results and Ablation Studies

Overall Results and Architecture Comparison Comparison with State-of-the-Art: Direct comparison with existing vein recognition methods is challenging due to divergent evaluation protocols. While most prior works report results under closed-set assumptions (e.g., [14, 20] reporting $>99\%$ accuracy), our evaluation enforces a stricter open-set protocol with unseen impostors. However, our method’s closed-set performance (Rank-1 $>99.6\%$ on MMCBNU_6000) is competitive with state-of-the-art closed-set systems, confirming that the open-set capability does not compromise identification accuracy for enrolled users.

We first compare different backbone architectures within our framework on the MMCBNU_6000 dataset. Table 2 reports the main open-set metrics (OSCR, AUROC, EER) together with closed-set Rank-1 accuracy and average inference time per query. This comparison allows us to select the best-performing architecture for subsequent experiments on cross-dataset generalisation and ablation studies.

We selected the MMCBNU_6000 dataset for backbone ablation and hyperparameter tuning primarily due to its scale and diversity. With 100 subjects, 600 distinct finger classes and 6,000 images, it provides a robust testbed for model selection.

Table 2. Architecture comparison on the MMCBNU_6000 dataset. We report OSCR, AUROC, EER, Rank-1 accuracy and average inference time per query.

Model	OSCR	AUROC	EER	R1	Accuracy	Time [ms]
ResNet50-CBAM	0.9945	0.9974	1.5714	0.9960	0.9833	1.7685
Attention U-Net	0.9894	0.9946	2.6825	0.9929	0.9627	2.4631
CNN	0.9867	0.9927	2.8452	0.9921	0.9746	1.5907
U-Net embedding	0.6314	0.7769	29.6071	0.7516	0.5317	2.6717

Based on these results, we select the ResNet50-CBAM backbone as our default model in the remaining experiments. Unless stated otherwise, all subsequent results and ablations are obtained with this backbone.

Cross-Dataset Generalisation and Domain Shift Analysis To assess the robustness of the selected model across acquisition setups, we evaluate it on all four datasets. Table 3 summarises the cross-dataset performance. While the framework maintains competitive open-set recognition on MMCBNU_6000 and FYO, we observe a marked performance drop on the UTFVP dataset (OSCR 0.76 vs 0.99 on MMCBNU).

Computational Analysis of Data Scarcity: This disparity reveals a critical dependency of Deep Metric Learning on the *scale of the training manifold*. MMCBNU provides 12,000 images for manifold approximation, whereas UTFVP offers only 1,440. Our analysis suggests that in strict open-set protocols, deep models require a critical mass of diverse samples to learn acquisition-agnostic features. The high EER on UTFVP (31.54%) indicates that without sufficient data density, the embedding space becomes sensitive to domain-specific artifacts (e.g., illumination variations), highlighting the necessity for transfer learning or synthetic data augmentation in small-scale medical scenarios.

Table 3. Cross-dataset performance of the proposed model on finger-vein (MMCBNU_6000, UTFVP), wrist-vein (FYO) and dorsal hand-vein datasets.

Dataset	OSCR	AUROC	EER	R1	Accuracy	Time [ms]
MMCBNU_6000	0.9945	0.9974	1.5714	0.9960	0.9833	1.7685
FYO	0.9746	0.9858	4.9424	0.9848	0.9627	2.4631
Dorsal	0.9226	0.9251	13.3513	0.9963	0.9746	1.5907
UTFVP	0.7576	0.7896	31.5418	0.9150	0.5317	2.6717

Ablation on Loss Functions We next investigate the impact of different loss functions on open-set performance. Table 4 reports the results for triplet loss, triplet loss with an additional centre regularisation term and contrastive loss, using the same backbone and training configuration. The comparison highlights which objective yields better separation between known and unknown score distributions.

Table 4. Ablation study on loss functions (MMCBNU_6000 dataset, fixed backbone and embedding dimension).

Loss	OSCR	AUROC	EER	R1	Accuracy	Time [ms]
Triplet	0.9925	0.9956	1.8968	0.9952	0.9833	1.7685
Triplet + Center	0.9920	0.9955	1.7619	0.9944	0.9627	2.4631
Contrastive	0.9770	0.9897	4.1865	0.9873	0.9746	1.5907

Ablation on Embedding Dimension Table 5 examines the effect of the embedding dimensionality on recognition and detection performance. We vary

the size of the embedding vector (128, 256, 384, 512) while keeping all other factors fixed. The results reveal the trade-off between representation capacity and compactness.

Table 5. Ablation study on embedding dimension (MMCBNU_6000 dataset, fixed backbone and loss).

Embedding dim	OSCR	AUROC	EER	R1	Accuracy	Time [ms]
128	0.9945	0.9974	1.5714	0.9960	0.9833	1.7685
256	0.9920	0.9955	1.7619	0.9944	0.9746	1.5907
384	0.9846	0.9930	3.2381	0.9889	0.5317	2.6717
512	0.9924	0.9966	1.5992	0.9944	0.9627	2.4631

Ablation on the Decision Rule Finally, we analyse the influence of the decision rule by comparing the standard 1-NN classifier with k -NN variants that average the top- k prototype similarities. Table 6 shows that modest values of k do not improve open-set performance in our setting: while Rank-1 accuracy remains high, OSCR and AUROC degrade and EER increases as k grows.

Table 6. Impact of the number of neighbours k in the decision rule (MMCBNU_6000 dataset).

k	OSCR	AUROC	EER	R1
1	0.9920	0.9955	1.7619	0.9944
3	0.8848	0.8889	18.8214	0.9937
5	0.7949	0.7991	27.3571	0.9937

Practical Deployment Considerations From a deployment perspective, the proposed prototype-based framework offers significant scalability advantages. Memory requirements for the embedding gallery are minimal: each enrolled class is represented by a single 512-dimensional vector (approximately 2 KB in 32-bit precision), meaning a database of one million identities requires only around 2 GB of memory. Furthermore, real-time enrolment of new users is seamless, requiring only a single forward pass to compute and store the new prototype without any network retraining. For highly scaled deployments, the exhaustive 1-NN search can be easily replaced with Approximate Nearest Neighbor (ANN) search algorithms to maintain real-time throughput.

5 Conclusion

We presented a computational framework for open-set vein biometric recognition, explicitly analyzing the boundaries of deep metric learning under strict

subject-disjoint protocols. By constraining the embedding manifold to a unit hypersphere and optimizing it via batch-hard triplet loss, our ResNet50-CBAM model achieves highly competitive open-set performance on large-scale benchmarks (OSCR 99.45% on MMCBNU_6000).

Crucially, our cross-dataset evaluation reveals a significant computational trade-off: while the method demonstrates high robustness on large datasets, performance degradation on smaller subsets (e.g., UTFVP) highlights the sensitivity of metric learning to the density of the training manifold. This confirms that while deep embeddings provide a solid foundation for secure biometrics, their deployment in data-scarce medical environments requires careful consideration of domain shift and sample diversity.

Future work will focus on addressing these computational constraints. First, to specifically mitigate the sensitivity to domain shift and data scarcity observed in our cross-dataset evaluations, we aim to investigate domain adaptation and generative synthetic data augmentation techniques. Second, we plan to explore uncertainty estimation methods to refine rejection thresholds dynamically. Third, given the privacy-sensitive nature of biometric data, we intend to integrate template protection mechanisms (e.g., cancelable biometrics) directly into the embedding pipeline. Finally, we will extend our evaluation to continual learning scenarios, enabling the system to adapt to new subjects without full retraining.

References

1. Abdullahi, S.B., et al.: Sequence-wise multimodal biometric fingerprint and finger-vein recognition network (stmfpfv-net). *Intelligent Systems with Applications* **19**, 200256 (2023). <https://doi.org/10.1016/j.iswa.2023.200256>
2. Al-Khafaji, R., Al-Tamimi, M.: Vein biometric recognition methods and systems: A review. *Advances in Science and Technology – Research Journal* **16**, 36–46 (2022). <https://doi.org/10.12913/22998624/144495>
3. Barcina-Blanco, M., et al.: Managing the unknown: a survey on open set recognition and tangential areas (2024)
4. Chen, Z., Yu, W., Bai, H., Li, Y.: An arcloss-based and openset-test-oriented finger vein recognition system. In: *Biometric Recognition*. pp. 287–294. Springer (2021)
5. Geng, C., Huang, S.J., Chen, S.: Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(10), 3614–3631 (2021). <https://doi.org/10.1109/tpami.2020.2981604>
6. Hemis, M., Kheddar, H., Bourouis, S., Saleem, N.: Deep learning techniques for hand vein biometrics: A comprehensive review (2024), <https://arxiv.org/abs/2409.07128>
7. Hou, B., Zhang, H., Yan, R.: Finger-vein biometric recognition: A review. *IEEE Transactions on Instrumentation and Measurement* **71**, 1–26 (2022). <https://doi.org/10.1109/TIM.2022.3200087>
8. Jain, A.K., Nandakumar, K., Ross, A.: 50 years of biometric research: Accomplishments, challenges, and opportunities. *Pattern Recognition Letters* **79**, 80–105 (2016). <https://doi.org/10.1016/j.patrec.2015.12.013>

9. Kocakulak, M., Avci, A., Acir, N.: A survey of finger-vein recognition using deep learning: Concepts, challenges, and opportunities. *Academic Platform Journal of Engineering and Smart Systems* **13**, 71–93 (2025). <https://doi.org/10.21541/apjess.1672743>
10. Lu, Y., Xie, S.J., Yoon, S., Yang, J.C., Park, D.S.: Robust finger vein roi localization based on flexible segmentation. *Sensors (Basel, Switzerland)* **13**, 14339 – 14366 (2013)
11. Lu, Y., Xie, S., Yoon, S., Wang, Z., Park, D.S.: An available database for the research of finger vein recognition. In: 2013 6th International Congress on Image and Signal Processing (CISP) (2014). <https://doi.org/10.1109/CISP.2013.6744030>
12. Mahdavi, A., Carvalho, M.: A survey on open set recognition. In: 2021 IEEE Fourth International Conference on Artificial Intelligence and Knowledge Engineering (AIKE). p. 37–44. IEEE (2021). <https://doi.org/10.1109/aike52691.2021.00013>
13. Marattukalam, F., Abdulla, W., Cole, D., Gulati, P.: Deep learning-based wrist vascular biometric recognition. *Sensors* **23**(6) (2023). <https://doi.org/10.3390/s23063132>
14. Radzi, S.A., Hani, M.K., Bakhteri, R.: Finger-vein biometric identification using convolutional neural network. *Turkish Journal of Electrical Engineering and Computer Sciences* **24**(3), 1863–1878 (2016). <https://doi.org/10.3906/elk-1311-43>
15. Ren, H., Sun, L., Guo, J., Han, C., Wu, F.: Finger vein recognition system with template protection based on convolutional neural network. *Knowledge-Based Systems* **227**, 107159 (2021). <https://doi.org/10.1016/j.knosys.2021.107159>
16. Singh, M., Singh, R., Ross, A.: A comprehensive overview of biometric recognition: approaches, challenges and future directions. *Journal of Ambient Intelligence and Humanized Computing* **12**, 1–23 (2021). <https://doi.org/10.1007/s12652-021-03123-w>
17. Ton, B.T., Veldhuis, R.: University of twente finger vascular pattern (utfvp). Dataset (2013), <https://www.utwente.nl/en/eemcs/dmb/downloads/utfvp/>
18. Toygar, , Babalola, F.O., Bitirim, Y.: Fyo: A novel multimodal vein database with palmar, dorsal and wrist biometrics. *IEEE Access* **8**, 82461–82470 (2020). <https://doi.org/10.1109/ACCESS.2020.2991475>
19. Zhang, J., Lu, Z., Li, M., Wu, H.: Gan-based image augmentation for finger-vein biometric recognition. *IEEE Access* **7**, 183118–183132 (2019). <https://doi.org/10.1109/ACCESS.2019.2960411>
20. Zhang, Y., Li, W., Zhang, L., Ning, X., Sun, L., Lu, Y.: Agcnn: Adaptive gabor convolutional neural networks with receptive fields for vein biometric recognition. *Concurrency and Computation: Practice and Experience* **34**(12), e5697 (2022). <https://doi.org/10.1002/cpe.5697>